Spatio-temporal Compressed Sensing with Coded Apertures and Keyed Exposures

Zachary T. Harmany, *Student Member, IEEE,* Roummel F. Marcia, *Member, IEEE,*
and Rebecca M. Willett, *Senior Member, IEEE*

arXiv:1111.7247v2 [stat.AP] 6 Jan 2012

***Abstract*—Optical systems which measure independent random projections of a scene according to compressed sensing (CS) theory face a myriad of practical challenges related to the size of the physical platform, photon efficiency, the need for high temporal resolution, and fast reconstruction in video settings. This paper describes a coded aperture and keyed exposure approach to compressive measurement in optical systems. The proposed projections satisfy the Restricted Isometry Property for sufficiently sparse scenes, and hence are compatible with theoretical guarantees on the video reconstruction quality. These concepts can be implemented in both space and time via either amplitude modulation or phase shifting, and this paper describes the relative merits of the two approaches in terms of theoretical performance, noise and hardware considerations, and experimental results. Fast numerical algorithms which account for the nonnegativity of the projections and temporal correlations in a video sequence are developed and applied to microscopy and short-wave infrared data.**

***Index Terms*—Compressive sensing, coded apertures, convex optimization, sparsity, coded exposure, Toeplitz matrices**

I. Introduction

THE theory of compressed sensing (CS) suggests that we can collect high-resolution imagery with relatively few photodetectors or a small focal plane array (FPA) when the scene is sparse or compressible in some dictionary or basis and the measurements are chosen appropriately [1], [2]. There has been significant recent interest in building imaging systems in a variety of contexts to exploit CS (*cf.* [3], [4], [5], [6], [7], [8], [9], [10], [11], [12]). By designing optical sensors to collect measurements of a scene according to CS theory, we can use sophisticated computational methods to infer critical scene structure and content. One particularly famous example in optical imaging is the Single Pixel Camera [13], which collects individual projections of the scene sequentially. While these measurements are supported by the CS literature, there are several practical challenges associated with the tradeoff between temporal resolution physical system footprint. In this paper we describe an alternative approach to designing a low frame-rate *snapshot* CS camera which naturally parallelizes the compressive data acquisition. Our approach is based on two imaging techniques called coded apertures and keyed exposures, which we explain next.

Coded apertures [14], [15] have a long history in low-light astronomical applications. Coded apertures were first developed to increase the amount of light hitting a detector in an optical system without sacrificing resolution (by, say, increasing the diameter of an opening in a pinhole camera). The basic idea is to use a mask, *i.e.,* an opaque rectangular plate with a specified pattern of openings, that allows significantly brighter observations with higher signal-to-noise ratio than those from conventional pinhole cameras [14], [15]. These masks encode the image before detection, inducing a more complicated point spread function than that associated with a pinhole aperture. The original scene is then recovered from the encoded observations in post-processing using an appropriate reconstruction algorithm which exploits the mask pattern. These multiplexing techniques are particularly popular in astronomical [16], [17] and medical [18], [19], [20] applications because of their efficacy at wavelengths where lenses cannot be used, but recent work has also demonstrated their utility for collecting both high resolution images and object depth information simultaneously [21].

Keyed exposure (also called coded exposure [22], flutter shutter [23], or coded strobing [24]) imaging was initially developed to facilitate motion deblurring in video using a relatively low frame rate. In some cases motion has been inferred from a single keyed exposure snapshot. The basic idea is that the camera sensor continuously collects light while the shutter is rapidly opened and closed; the shutter movement effectively modulates the motion blur point spread function, and with well-chosen shutter movement patterns it becomes possible to deblur moving objects. Similar effects can be achieved using a strobe light instead of moving a shutter.

Despite the utility of the above methods in specific settings, they both face some limitations. The design of conventional coded apertures does not account for the inherent structure and compressibility of natural scenes, nor the potential for nonlinear reconstruction algorithms. Likewise, existing keyed exposure methods focus on direct (uncoded) measurements of the spatial content of the scene and have limited reconstruction capabilities, as we detail below. The Coded Aperture Keyed Exposure (CAKE) sensing paradigm we propose in this paper is designed to allow nonlinear high-resolution video reconstruction from relatively few measurements in more general settings.

This research is supported by DARPA Contract No. HR0011-04-C-0111, ONR Grant No. N00014-06-1-0610, DARPA Contract No. HR0011-06-C-0109, and NSF-DMS-08-11062. Portions of this work were presented at the IEEE International Conference on Acoustic, Speech, and Signal Processing, March 2008, and at the SPIE Electronic Imaging Conference, January 2009, and SPIE Photonics Europe, April 2010.

Z.T. Harmany and R.M. Willett are with the Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708 USA (e-mail: zth@duke.edu, willett@duke.edu).

R.F. Marcia is with the School of Natural Sciences, University of California, Merced, CA 95343 USA (e-mail: rmarcia@ucmerced.edu).

A. Problem Formulation

We consider the problem of reconstructing an N-frame video sequence $f^\star$, where each frame is an $n_1 \times n_2$ two-dimensional image denoted $f_t^\star$. Using standard vector representation, we have that $f_t^\star \in \mathbb{R}^n$ for $t = 1, \ldots, N$ where $n \triangleq n_1 n_2$ is the total number of pixels. As a result, the vector representation of the video sequence is $f^\star = (f_1^\star, \ldots, f_N^\star) \in \mathbb{R}^{nN}$.

The observations y of $f^\star$ are also acquired as a video sequence. We do not assume that the observations are acquired at the same rate at which we will ultimately reconstruct $f^\star$. In general, we assume y is an M-frame video sequence, with each frame y_k of size $m_1 \times m_2$. Similarly to $f^\star$, we have $y_k \in \mathbb{R}^m$ for $k = 1, \ldots, M$, where $m \triangleq m_1 m_2$, therefore $y = (y_1, \ldots, y_M) \in \mathbb{R}^{mM}$.

We observe $f^\star$ via a spatio-temporal *sensing matrix* $A \in \mathbb{R}^{mM \times nN}$ which linearly projects the spatio-temporal scene onto an mM-dimensional set of observations:

$$y = Af^\star + w, \tag{1}$$

where $w \in \mathbb{R}^{mM}$ is noise associated with the physics of the sensor.

CS optical imaging systems must be designed to meet several competing objectives:

- The sensing matrix A must satisfy some necessary criterion (such as the RIP, defined below) which provides theoretical guarantees on the accuracy with which we can estimate $f^\star$ from y.
- The total number of measurements, mM, must be lower than the total number of pixels to be reconstructed, nN. This is achievable via compressive spatial acquisition ($m < n$), frame rate reduction ($M < N$), or simultaneous spatio-temporal compression.
- The measurements y must be *causally* related to the temporal scene $f^\star$, which restricts the structure of the projections A.
- The optical measurements modeled by A must be implementable in a way that results in a smaller, cheaper, more robust, or lower power system.
- The sensing matrix structure must facilitate fast reconstruction algorithms.

This paper demonstrates that compressive Coded Aperture Keyed Exposure systems achieve all these objectives.

B. Contributions

The primary contribution of this paper is the design and theoretical characterization of compressive Coded Aperture Keyed Exposure (CAKE) sensing. We explore amplitude modulating and phase shifting masks and describe theoretical and implementation aspects of both. We further describe how keyed exposure ideas can be used in conjunction with coded apertures to increase both the spatial and temporal resolution of video from relatively few measurements. We prove hitherto unknown theoretical properties of such systems and demonstrate their efficacy in several simulations. In addition, we discuss several important algorithmic aspects of our approach, including a mean-subtraction pre-processing step which allows us to sidestep challenging theoretical aspects associated with nonnegative sensing matrices (such as amplitude modulating coded apertures). This paper builds substantially upon earlier preliminary studies by the authors [25], [26], [27] and related independent work by Romberg [28].

C. Organization of the Paper

The paper is organized as follows. Section II-A describes conventional coded aperture imaging techniques and Section II-B describes keyed exposure techniques currently in the literature. We describe the compressive sensing problem and formulate it mathematically in Section II-C. In Section III, we show how CS theory can be used for constructing coded aperture masks that can easily be implemented for improving image reconstruction resolution in snapshot imaging; this includes theoretical results, a discussion of implementation details and tradeoffs in practical optical systems, and experimental results. We consider applications to video compressed sensing using the full CAKE paradigm in Section IV, including theory and experimental results.

II. Background

Prior to detailing our main contributions, we first review pertinent background material. This review touches upon the development of coded aperture imaging, coded exposure photography, and a brief review of salient concepts in compressed sensing theory.

A. Coded Aperture Imaging

Seminal work in coded aperture imaging includes the development of masks based on Hadamard transform optics [29] and pseudorandom phase masks [30]. Modified Uniformly Redundant Arrays (MURAs) [31] are generally accepted as optimal mask patterns for coded aperture imaging. These mask patterns (which we denote by h^{MURA}) are binary, square patterns, whose *grid size matches the spatial resolution of the photo-detector*. Each mask pattern is specifically designed to have a complementary pattern $\overline{h}^{\mathrm{MURA}}$ such that $h^{\mathrm{MURA}} * \overline{h}^{\mathrm{MURA}}$ is a single peak with flat side-lobes (*i.e.,* a Kronecker δ function).

In practice, the resolution of a detector array dictates the properties of the mask pattern and hence resolution at which $f^\star$ can be reconstructed. We model this effect as $f^\star$ being downsampled to the resolution of the detector array and then convolved with the mask pattern h^{MURA}, which has the same resolution as the FPA and the downsampled $f^\star$, *i.e.,*

$$y = (D_{\text{int}} f^\star) * h^{\mathrm{MURA}} + w, \tag{2}$$

where $*$ denotes circular convolution, w corresponds to noise associated with the physics of the sensor, and $D_{\text{int}} f^\star$ is the *integration* downsampling of the scene, which consists of partitioning $f^\star$ into uniformly sized $d_1 \times d_2$ blocks, where $d_i \triangleq n_i / m_i$ for $i = 1, 2$, and measuring the total intensity in each block.

Because of the construction of h^{MURA} and $\overline{h}^{\mathrm{MURA}}$, $D_{\mathrm{int}}f^\star$ can be reconstructed using

$$\widehat{f} = y * \overline{h}^{\mathrm{MURA}}.$$

However, the resulting resolution is often lower than what is necessary to capture some of the desired details in the image. Clearly, the estimates from MURA reconstruction are limited by the spatial resolution of the photo-detector. Thus, high resolution reconstructions cannot generally be obtained from low-resolution MURA-coded observations.

It can be shown that this mask design and reconstruction result in minimal reconstruction errors *at the FPA resolution* and *subject to the constraint that linear, convolution-based reconstruction methods would be used*. However, when the scene of interest is sparse or compressible, and nonlinear sparse reconstruction methods may be employed, then CS ideas can be used to design coded aperture which yield higher resolution images. Before describing the details of this, we briefly review two key relevant concepts from CS.

B. Coded (Keyed) Exposure Imaging

Coded (or keyed) exposures were developed recently in the computational photography community. Initial work in this area was focused on engineering the temporal component of a motion blur point spread function by rapidly opening and closing the shutter during a single exposure or a small number of exposures at a low frame rate [22], [23]. That is,

$$y = A^{\mathrm{KE}}f^\star + w = \sum_{i\in S} f_i^\star + w,$$

where the keyed exposure (KE) measurement matrix A^{KE} selects the subset of frames during which the shutter is open. We refer to this subset as the exposure code $S \subseteq \{1,\dots,N\}$.

If an object is moving during image acquisition, then a static shutter would induce a typical motion blur, making the moving object difficult to resolve with standard deblurring methods. However, by "fluttering" the shutter during the exposure using carefully designed patterns, the induced motion blur can be made invertible and moving objects can be accurately reconstructed. Instead of a moving shutter, more recent work uses a strobe light to produce a similar effect [24].

While this novel approach to video acquisition can produce very accurate deblurred images of moving objects, there is significant overhead associated with the reconstruction process. To see why, note that every object moving with a different velocity or trajectory will produce a different motion blur. This means that (a) any stationary background must be removed during preprocessing and (b) multiple moving objects must be separated and processed individually.

More recently, it was shown that these challenges could be sidestepped when the video is temporally periodic (*e.g.,* consider a video of an electronic toothbrush spinning) [24]. The periodic assumption amounts to a sparse temporal Fourier transform of the video, and this approach, called coded strobing, is a compressive acquisition in the temporal domain. As a result, the authors were able to leverage ideas from compressed sensing to achieve high-quality video reconstruction.

The assumption of a periodic video makes it possible to apply much more general reconstruction algorithms that do not require background subtraction or separating different moving components. However, it is a very strong assumption, which places some limits in its applicability to real-world settings. The approach described in this paper has similar performance guarantees but operates on much more general video sequences.

C. Compressed Sensing

In this section we briefly define the Restricted Isometry Property (RIP) and explain its significance to reconstruction performance. In subsequent sections, we demonstrate our primary theoretical contribution, which is to prove the RIP for compressive coded aperture and keyed exposure systems.

Definition 1 (Restricted Isometry Property (RIP) [32])**.** *A matrix A satisfies the RIP of order s if there exists a constant $\delta_s \in (0,1)$ for which*

$$(1-\delta_s)\|f\|_2^2 \le \|Af\|_2^2 \le (1+\delta_s)\|f\|_2^2. \tag{3}$$

holds for all s-sparse $f \in \mathbb{R}^n$. If this property holds, we say A is RIP(s,δ_s).

Matrices which satisfy the RIP are called CS matrices; and when combined with sparse recovery algorithms, they are guaranteed to yield accurate estimates of the underlying function $f^\star$:

Theorem 1 (Sparse Recovery with RIP Matrices [33], [34])**.** *Let A be a matrix satisfying RIP$(2s,\delta_{2s})$ with $\delta_{2s} < \sqrt{2}-1$, and let $y = Af + w$ be a vector of noisy observations of any signal $f \in \mathbb{R}^n$, where the w is a noise or error term with $\|w\|_2 \le \epsilon$ for some $\epsilon > 0$. Let f_s be the best s-sparse approximation of f; that is, f_s is the approximation obtained by keeping the s largest entries of f and setting the others to zero. Then the estimate*

$$\begin{aligned} \widehat{f} \;=\; & \underset{f\in\mathbb{R}^n}{\arg\min} && \|f\|_1 \\ & \text{subject to} && \|y - Af\|_2 \le \epsilon, \end{aligned} \tag{4}$$

obeys

$$\|f - \widehat{f}\|_2 \le C_{1,s}\epsilon + C_{1,s}\frac{\|f-f_s\|_1}{\sqrt{s}},$$

where $C_{1,s}$ and $C_{2,s}$ are constants which depend on s but not on n or m.

Note that the reconstruction (4) in Theorem 1 is equivalent to

$$\widehat{f} = \underset{f\in\mathbb{R}^n}{\arg\min}\, \tfrac{1}{2}\|y - Af\|_2^2 + \tau\|f\|_1 \tag{5}$$

where $\tau > 0$, which depends on ϵ, can be viewed as a regularization parameter.

III. Compressive Coded Apertures for Snapshot Imaging

This section first considers a snapshot acquisition model. Our goal is to recover a static high-resolution scene from a single image where all pixels are collected simultaneously. In

terms of the notation in Sec. I-A, we have $N = M = 1$, so that A is of size $m \times n$. The sensing matrix for compressive coded aperture (CCA) systems can be modeled mathematically as

$$Af^\star = D(f^\star * h); \tag{6}$$

where h is a coding mask, and D is a subsampling operator (detailed below). Here the coding mask, h, is at the size and resolution at which $f^\star$ will be reconstructed; this is in contradistinction to the MURA system, in which h^{MURA} is at the size and resolution of the FPA. Thus in (6), we model the measurements as the scene being convolved with the coded mask and *then* downsampled.

Using a similar model, Romberg [28] conducted related work concurrent and independent of our initial investigations [25]. We will summarize the key features of this model and compare with our approach in Sec. III-B. While these models share common elements, we will see in Sec. III-C that there are important tradeoffs associated with the theory and implementation of each strategy.

Recent work by Bajwa et al. [35], [36], showed that random circulant matrices (and Toeplitz matrices, in general) are sufficient to recover sparse $f^\star$ from y with high probability. In particular, they showed that a Toeplitz matrix whose first row contains elements drawn independently from a Gaussian distribution are RIP(s, δ_s) when $m \geq Cs^2 \log(n)$ for some constant C. Here, we extend these results to pseudo-circulant matrices and use them to motivate our mask design. Our model differs from that in [28] in that we consider a different generative model for the coded aperture mask, as well as a different subsampling strategy:

- The elements of the coding mask h are generated iid according to a particular generating distribution (*e.g.,* an appropriately scaled Rademacher, uniform, or Gaussian distribution).
- Our analysis allows for deterministic downsampling in which we collect one sample per nonoverlapping block.

In our notation, we distinguish mask patterns in this manner using the abbreviations BS, US, and GS, where the first character denotes the distribution used to generate the mask, *i.e.,* binary, uniform, or Gaussian, and the second is a reminder that these masks are generated directly in the spatial domain.

A. CCAs Generated in the Spatial Domain

The two-dimensional convolution of h with an image $f^\star$ as in (6) can be represented as the application of the Fourier transform to $f^\star$ and h, followed by element-wise multiplication and application of the inverse Fourier transform. In matrix notation, this series of linear operations can be expressed as

$$\text{vect}(f^\star * h) = \mathcal{F}^{-1} \operatorname{diag}(\mathcal{F}h)\mathcal{F}f^\star = Rf^\star, \tag{7}$$

where $\text{vect}(f)$ is a vectorized representation of an image f, $\mathcal{F}$ is the two-dimensional Fourier transform matrix, and $\operatorname{diag}(\mathcal{F}h)$ is a diagonal matrix whose elements correspond to the transfer function, which is the Fourier transform of h. The matrix product $R = \mathcal{F}^{-1} \operatorname{diag}(\mathcal{F}h)\mathcal{F} \in \mathbb{R}^{n \times n}$ is block-circulant and each block is in turn circulant. In matrix notation,

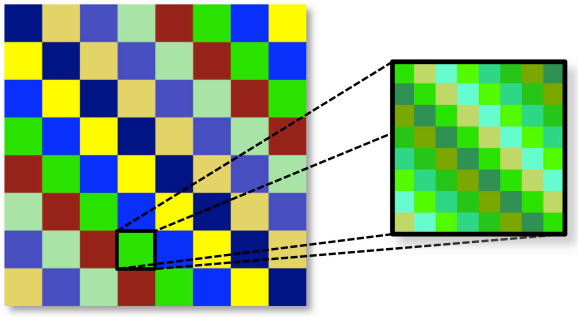

Fig. 1. The $n \times n$ matrix $\mathcal{F}^{-1} \operatorname{diag}(\mathcal{F}h)\mathcal{F}$ is block-circulant with n_2 blocks in each row and column. Each block is $n_1 \times n_1$ and is circulant.

R is consists of $n_2 \times n_2$ blocks,

$$R = \begin{bmatrix} R_1 & R_{n_2} & \cdots & R_3 & R_2 \\ R_2 & R_1 & \cdots & R_4 & R_3 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ R_{n_2} & R_{n_2-1} & \cdots & R_2 & R_1 \end{bmatrix}, \tag{8}$$

where each $R_j \in \mathbb{R}^{n_1 \times n_1}$ is circulant; *i.e.,* R_j is of the form

$$R_j = \begin{bmatrix} r_{j,1} & r_{j,n_1} & \cdots & r_{j,3} & r_{j,2} \\ r_{j,2} & r_{j,1} & \cdots & r_{j,4} & r_{j,3} \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ r_{j,n_1} & r_{j,n_1-1} & \cdots & r_{j,2} & r_{j,1} \end{bmatrix},$$

(see Fig. 1). This block-circulant with circulant-block (BCCB) structure of R is a direct result of the fact that the k-point one-dimensional Fourier transform F_k diagonalizes any $k \times k$ circulant matrix (such as R_j with $k = n_1$) and so $\mathcal{F} \equiv F_{n_2} \otimes F_{n_1}$ diagonalizes block-circulant matrices (such as R). Here $\otimes$ denotes the Kronecker matrix product.

We now examine the generation of a compressed sensing matrix from the convolution R by restricting the number of measurements collected of the vector $Rf^\star$. Here we simply subsample the vector $Rf^\star$ by applying a pointwise subsampling matrix D_{sub}. The operation of applying D_{sub} consists of retaining only one measurement per uniformly sized $d_1 \times d_2$ block, *i.e.,* we subsample by d_1 in the first coordinate and d_2 in the second coordinate so that the result is an $m_1 \times m_2$ image with $m_i = n_i/d_i, i = 1, 2$. In matrix form D_{sub} can be thought of as retaining a certain number of rows of the identity matrix. Because of the structure and deterministic nature of this type of downsampling, it is often more straightforward to realize in practical imaging hardware (see Sec. III-C).

The resulting projection matrix A is then given by

$$A = D_{\text{sub}}R = D_{\text{sub}}\mathcal{F}^{-1} \operatorname{diag}(\mathcal{F}h)\mathcal{F}.$$

We now show that if the elements of h are drawn from an appropriate probability distribution, then A will be RIP(s, δ_s) with high probability.

Theorem 2 (Spatial-Domain CCA Sensing)**.** *Let h^{BS} be a mask with entries generated i.i.d. according to the scaled*

Rademacher distribution

$$h^{\mathrm{BS}}_{k_1,k_2} = \begin{cases} \sqrt{d/n} & \textit{with probability } 1/2, \\ -\sqrt{d/n} & \textit{with probability } 1/2, \end{cases} \tag{9}$$

for $k_1 = 1, \ldots, n_1, k_2 = 1, \ldots, n_2$. Let $A^{\mathrm{BS}} = D_{sub}R$ be an $m \times n$ matrix, where $R \in \mathbb{R}^{n\times n}$ is the BCCB matrix generated from h^{BS}, and $D_{sub} \in \mathbb{R}^{m\times n}$ is the pointwise subsampling matrix. Then, there exists constants $c_1, c_2 \geq 0$ depending on δ_s such that for any

$$m \geq c_1 s^2 \log(n), \tag{10}$$

A^{BS} is RIP(s, δ_s) with probability exceeding

$$1 - 2n^2 e^{-c_2 m/s^2}. \tag{11}$$

Remark 1. *This binary-valued distribution was selected to model coded apertures with two states – open and closed – per mask element. We discuss the issue of implementing these masks in Section III-C1.*

Remark 2. *It is straightforward to extend the result to other mask generating distributions, such as uniform and Gaussian, which we denote h^{US} and h^{GS}, with associated sensing matrices A^{US} and A^{GS}.*

We present the proof of Theorem 2 in Appendix A.

For successful recovery, the coded aperture masks are designed to be satisfy RIP$(2s, \delta_{2s})$ as described in (3) with high probability when $m \geq c_1 s^2 \log(n)$. Note that this is somewhat weaker than what can be achieved by a purely unstructured i.i.d. randomly generated sensing matrix which satisfies (3) with high probability when $m \geq \tilde{c}_1 s \log(n/s)$ for some constant $\tilde{c}_1 > 0$. Intuition may lead one to believe the extra factor of s is due to the fact that the m projections sensed using the amplitude modulated mask framework exhibit dependencies. However, in many settings this theoretical disadvantage is offset by advantageous practical implementations.

Sparse Gradients Scenes: The above theory is applicable to reconstructing $f^\star$ when $f^\star$ is sparse in the pixel basis; similar results hold when the *gradient* of $f^\star$ is sparse. For simplicity of notation we will show then in a 1d setting; the extension to 2d is straightforward. Let ∇f denote the first-order gradient of f, so that

$$\nabla \triangleq \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 \\ -1 & 1 & 0 & \ddots & \vdots \\ 0 & -1 & 1 & \ddots & 0 \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & -1 & 1 \end{bmatrix}. \tag{12}$$

As noted in [35], $f^\star = \nabla^{-1}\theta$, where θ is the sparse gradient image. Thus if we sense $y = A\nabla f^\star + w = A\nabla\nabla^{-1}\theta + w = A\theta + w$, we can expect to recover θ using sparse reconstruction methods.

B. CCAs Generated in the Fourier Domain

The random convolution in [28] follows the same structure as that in Sec. III-A. However, in that work the convolution is generated randomly in the frequency domain. More specifically, the entries of the transfer function correspond to random *phase shifts* (with some constraints to keep the resulting observation matrix real-valued). We denote the resulting convolution kernel as h^{UP}, where "UP" refers to "uniform phase". For simplicity of presentation, we describe the generating distribution for a one-dimensional convolution for even-length masks. In particular, let $\sigma = Fh^{\mathrm{UP}}$, and generate σ such that for $k = 1, \ldots, n$,

$$\sigma_k = \begin{cases} \pm 1 \text{ with equal probability} & \text{if } k = 1, \\ e^{i\phi} \text{ with } \phi \sim \mathcal{U}(0, 2\pi) & \text{if } 2 \leq k \leq n/2, \\ \pm 1 \text{ with equal probability} & \text{if } k = n/2 + 1 \\ \sigma^*_{n-k+2} & \text{if } n/2 + 2 \leq k \leq n. \end{cases} \tag{13}$$

Here $\mathcal{U}(0, 2\pi)$ denotes the uniform distribution over $[0, 2\pi]$. The real-valued convolution kernel is then given by $h^{\mathrm{UP}} = F^{-1}\sigma$. This result can be extended easily for two-dimensional convolutions.

To form a compressed sensing matrix from this random convolution, [28] considers two different downsampling strategies: sampling at random locations, and random demodulation. The first method entails selecting a random subset $\Omega \subset \{1, \ldots, n\}$ of indices. We form a downsampling matrix D_Ω by retaining only the rows of the identity matrix indexed by Ω, hence the resulting measurement matrix is given by

$$A^{\mathrm{UP}} = D_\Omega \mathcal{F}^{-1} \operatorname{diag}(\sigma) \mathcal{F}, \tag{14}$$

where σ is generated according to the two-dimensional analog of (13). The random demodulation method multiplies the result of the random convolution by a random sign sequence $s \in \{-1, 1\}^n$, such that $s_i = \pm 1$ with equal probability for all $i = 1, \ldots, n$, then performs an integration downsampling of the result. Therefore in this case the measurement matrix is

$$A^{\mathrm{UP}} = D_{\mathrm{int}} \operatorname{diag}(s) \mathcal{F}^{-1} \operatorname{diag}(\sigma) \mathcal{F}. \tag{15}$$

It can be shown that both of these strategies yield RIP-satisfying matrices with high probability.

Theorem 3 (Fourier-Domain CCA Sensing)**.** *Let A^{UP} be an $m \times n$ sensing matrix resulting from random convolution with phase shifts followed by a downsampling strategy, and let W denote an arbitrary orthonormal basis. If the downsampling is random subsampling as in* (14)*, then there exists constant $c_3 > 0$ such that with probability exceeding $1 - O(n^{-1})$, for*

$$m \geq c_3 \delta^{-2} \min(s \log^6 n, s^2 \log^2 n),$$

$A^{\mathrm{UP}}W$ satisfies RIP$(2s, \delta_{2s})$ with $\delta_{2s} \leq \delta$. If the downsampling is random demodulation as in (15)*, then there exists constant $c_4 > 0$ such that with probability exceeding $1 - O(n^{-1})$, for*

$$m \geq c_4 \delta^{-2} \min(s \log^6 n, s^2 \log n),$$

$A^{\mathrm{UP}}W$ satisfies RIP$(2s, \delta_{2s})$ with $\delta_{2s} \leq \delta$.

These theoretical results are stronger than those in Theorem 2, especially since they allow sparsity in arbitrary orthonormal bases. However, there are important differences between the observation models which have a significant impact on the feasibility and ease of hardware implementation. We elaborate on this in Section III-C. Nevertheless, the theory developed in [28] lends important theoretical support to the general concept of compressive coded apertures.

C. Hardware and Practical Implementation Considerations

In this section we describe how we shift from modeling the coded aperture masks in a way that is compatible with compressed sensing theory, to a model that describes their actual implementation in an optical system. Our analysis does not account for the bandwidth of the lenses; in particular, we implicitly assume that the bandwidth of the lenses is high enough that band-limitation effects are negligible at the resolution of interest. In all of the hardware settings described below, precise alignment of the optical components (*e.g.,* the mask and the focal plane array) is critical to the performance of the proposed system. Often a high-resolution FPA is helpful for alignment and calibration.

In this paper we focus on *incoherent* light settings (consistent with many applications in astronomy, microscopy, and infrared imaging). In this case, the coded aperture must be real-valued and flux-preserving (*i.e.,* the light intensity hitting the detection cannot exceed the light intensity of the source). In this section, we consider the following apertures:

- **Binary Spatial Mask:** $h \in \{0, 1/n\}^{n_1 \times n_2}$, drawn with equal probability,
- **Uniform Spatial Mask:** $h \in [0, 1/n]^{n_1 \times n_2}$, where each element is drawn independently from a uniform distribution, or
- **Uniform Phase Mask:** $h = |\mathcal{F}p|^2$ for some p, where p corresponds to a phase-shifting mask in an incoherent light setting.

1) Amplitude Modulation Masks: In a conventional lensless coded aperture imaging setup, the point spread function associated with the aperture is the mask pattern h itself. To shift a RIP-satisfying aperture as in Theorem 2 to an implementable aperture, one simply needs to apply an affine transform to h mapping $[-\sqrt{d/n}, \sqrt{d/n}]$ to $[0, 1/n]$. This transform ensures that the resulting mask pattern is nonnegative and flux-preserving.

These amplitude modulating masks may be implemented using a spatial light modulator (SLM) or placing chrome on quartz. The SLMs may be preferable in video settings where the underlying scene contains motion and using a different mask pattern at each time step boosts performance. However, in order for the proposed approach to work, the mask or SLM used must be higher resolution than the FPA. Currently, very high resolution SLMs are still in development. Chrome on quartz masks can be made with higher resolution than many SLMs, but cannot be changed on the fly unless we mount a small number of fixed masks on a rotating wheel or translating stage. The uniform amplitude modulation masks in particular could be constructed using a high-resolution halftoning procedure, which is easiest to implement at the necessary resolution with chrome on quartz.

Both Robucci *et al.* [37] and Majidzadeh *et al.* [38] have proposed performing the analog, random convolution step in complementary, metal-oxide-semiconductor (CMOS) electronics. A clear advantage to this architecture is that the additional optics required for spatial light modulation are removed in favor of additional circuitry, immediately reducting imager size.

2) Phase Shift Masks: Phase shifting masks for coded aperture imaging have been implemented recently using a phase screen [39]. This approach allows one to account for diffraction in the optical design. However, depending on the precise optical architecture, phase shift masks may be much less photon-efficient than amplitude modulation masks. Additionally, the mask generation distribution described in Eq. (13) will result in negative entries for the corresponding PSF $h^{\mathrm{UP}} = \mathcal{F}^{-1}\sigma$. To compensate for this, the phase mask must be mean-shifted to make all entries nonnegative, so that we are actually implementing

$$h_+^{\mathrm{UP}} = c(h^{\mathrm{UP}} - \min(h^{\mathrm{UP}})\mathbf{1}), \tag{16}$$

with the constant c selected so that the implementable PSF h_+^{UP} is flux-preserving and $\mathbf{1}$ is a matrix of ones and of the same dimension as h^{UP}.

3) Implementable Masks for Scenes with Sparse Gradients: As described in Section III-A, theoretical results for scenes with sparse gradients hold for the sparse difference operator ∇ defined in (12). The problem with this approach is that ∇ is invertible but not circulant – and hence the sensing matrix $A\nabla$ cannot be implemented with a coded aperture system. We address this by noting that if the upper right element of ∇ were -1 instead of 0, then the resulting sensing matrix, denoted $\widetilde{\nabla}$, could be implemented physically and is a close approximation to the theoretically supported ∇. In short, the theoretically supported solution is to set

$$\begin{aligned} y &= A\nabla f^\star + w \\ \widehat{\theta} &= \arg\min_\theta \tfrac{1}{2}\|y - A\theta\|_2^2 + \tau\|\theta\|_1 \\ \widehat{f} &= \nabla^{-1}\widehat{\theta} \end{aligned}$$

while an implementable close approximation is to set $A^{\mathrm{G}} \triangleq A\widetilde{\nabla}$ and

$$\begin{aligned} y &= A^{\mathrm{G}} f^\star + w \\ \widehat{f} &= \arg\min_f \tfrac{1}{2}\|y - A^{\mathrm{G}} f\|_2^2 + \tau\|f\|_{\mathrm{TV}}, \end{aligned}$$

where $\|\cdot\|_{\mathrm{TV}}$ is the total variation seminorm which causes $\widehat{f}$ to have sparse gradients. In our numerical experiments we compare the performance of TV regularization to sparsity penalization for the task of reconstructing static scenes.

4) Downsampling Implementation: In developing RIP-satisfying AMM coded apertures, Theorem 2 assumes the subsampling operation selects one measurement per $d_1 \times d_2$ block. From an implementation standpoint, this operation effectively discards a large portion $((d-1)/d)$ of the available light, which would result in much lower signal-to-noise ratios

at the detector. A more pragmatic approach is to use larger detector elements that essentially sum the intensity over each $d_1 \times d_2$ block, making a better use of the available light. We call this operation *integration downsampling* to distinguish it from subsampling. The drawback to this approach is that we lose many of the desirable features of the system in terms of the RIP. Integration downsampling causes a large coherence between neighboring columns of the resulting sensing matrix A. An intermediate approach would randomly sum a fraction of the elements in each size d block, which increases the signal-to-noise ratio versus subsampling, but yields smaller expected coherence. This approach is motivated by the random demodulation proposed in [28] and described in Sec III-B, whereby the signal is multiplied by a random sequence of signs $\{-1, +1\}$, then block-wise averaged. The pseudo-random summation proposed here can be thought of as an optically realizable instantiation of the same idea where we multiply by a random binary $\{0, 1\}$ sequence. We explore the effects of these choices in our numerical results section.

5) Noise and Quantization: While CS is particularly useful when the FPA needs to be kept compact, it should be noted that CS is more sensitive to measurement errors and noise than more direct imaging techniques. The experiments conducted in this paper simulated very high signal-to-noise ratio (SNR) settings and showed that CS methods can help resolve high resolution features in images. However, in low SNR settings CS reconstructions can exhibit significant artifacts that may even cause more distortion than the low-resolution effects associated with conventional coded aperture techniques such as MURA.

Similar observations are made in [40], which presents a direct comparison of the noise robustness of CS in contrast to conventional imaging techniques both in terms of bounds on how reconstruction error decays with the number of measurements and in a simulation setup; the authors conclude that for most real-world images, CS yields the biggest gains in high signal-to-noise ratio (SNR) settings. Related theoretical work in [41] show that in the presence of low SNR photon noise, theoretical error bounds can be large, and thus the expected performance of CS may be limited unless the number of available photons to sense is sufficiently high. These considerations play an important role in choosing the type of downsampling to implement.

Similar issues arise when considering the bit-depth of focal plane arrays, which corresponds to measurement quantization errors. Future efforts in designing optical CS systems must carefully consider the amount of noise anticipated in the measurements to find the optimal tradeoff between the focal plane array size and image quality.

D. Reconstruction

To solve the CS minimization problem (5), we use well-established gradient-based optimization methods. They are particularly suitable in our setting because the block-circulant structure of A described in the previous section allows for very fast matrix-vector products that are critical to the speed of their performance. However, our observation matix A is not zero mean and, therefore, negatively impacts the performance of these CS-based reconstruction algorithms. In this section, we describe how we take full advantage of the structure of A and address how we mitigate the negative effects of A having nonzero-mean.

1) Sparsity-promoting Methods: We note that most popular and effective methods for performing the above reconstruction are iterative algorithms with repeated applications of the operators A and $(A)^T$ to scene estimates and residuals [42], [43]. In most CS settings, computing these matrix-vector multiplications is a large computational burden. However, because of the circulant structure of A, this computation is very efficient using fast Fourier transform algorithms. In particular, we compared the computation time for reconstructing a 256×256 scene using 128×128 CCA measurements with the time required to reconstruct the same scene using the same number of measurements computing using a dense random projection matrix; our reconstruction was 250 times faster because of our exploitation of the Toeplitz structure in A.

2) Mean Subtraction: Generative models for random projection matrices used in CS involve drawing elements independently from a zero-mean probability distribution [1], [44], [32], [35], [45], and likewise a zero-mean distribution was used to analyze the coded aperture masks described in Sec. III. However, as coded aperture masks with zero mean are not physically realizable in optical systems, we generate our physically realizable masks from an appropriately scaled Bernoulli distribution with values $\{0, 1/n\}$. This shifting ensures that the coded aperture corresponds to an implementable (*i.e.,* nonnegative and intensity preserving) mask pattern which does satisfies the assumptions needed to model photon propagation through an optical system.

While necessary to accurately model real-world optical systems, this shifting negatively impacts the performance of well-established ℓ_2-ℓ_1 reconstruction algorithms for the following reason. Several of these algorithms (*e.g.,* [42], [43]) assume that the ℓ_2 data fidelity term $\phi(f) = \frac{1}{2}\|Af - y\|_2^2$ (proportional to the negative log Gaussian likelihood) is such that $\nabla^2\phi(x) = A^T A$ can be well-approximated by αI, $\alpha > 0$. This assumption is crucial to the performance of these algorithms.

If we collect measurements using a realizable mask with elements in the range $[0, 1/n]$, the resulting matrix A does not satisfy this condition due to the mean offset. Therefore in the reconstruction we use a mean-shifted sensing matrix

$$A^0 = A - \mu_A \mathbf{1}_{m\times n},$$

where $\mu_A = (1/mn)\mathbf{1}_m^T A \mathbf{1}_n$ is the mean value of A. This new matrix is such that $(A^0)^T A^0 \approx \alpha I$ is a valid assumption. However, to use this matrix in the reconstruction, we need to adjust our estimate accordingly. To compensate for this, we decompose our estimate f into its mean component μ_f and a zero-mean deviation f^0: $f = f^0 + \mu_f \mathbf{1}_n$. So then

$$\begin{aligned} Af &= (A^0 + \mu_A \mathbf{1}_{m\times n})(f^0 + \mu_f \mathbf{1}_n) \\ &= A^0 f^0 + n\mu_A\mu_f \mathbf{1}_m + \mu_f A^0 \mathbf{1}_n + \mu_A(\mathbf{1}_n^T f^0)\mathbf{1}_m \\ &\approx A^0 f^0 + n\mu_A\mu_f \mathbf{1}_m. \end{aligned}$$

The data y can be decomposed in a similar fashion, so that $y = y^0 + \mu_y \mathbf{1}_m$. Thus a straightforward estimate for the mean

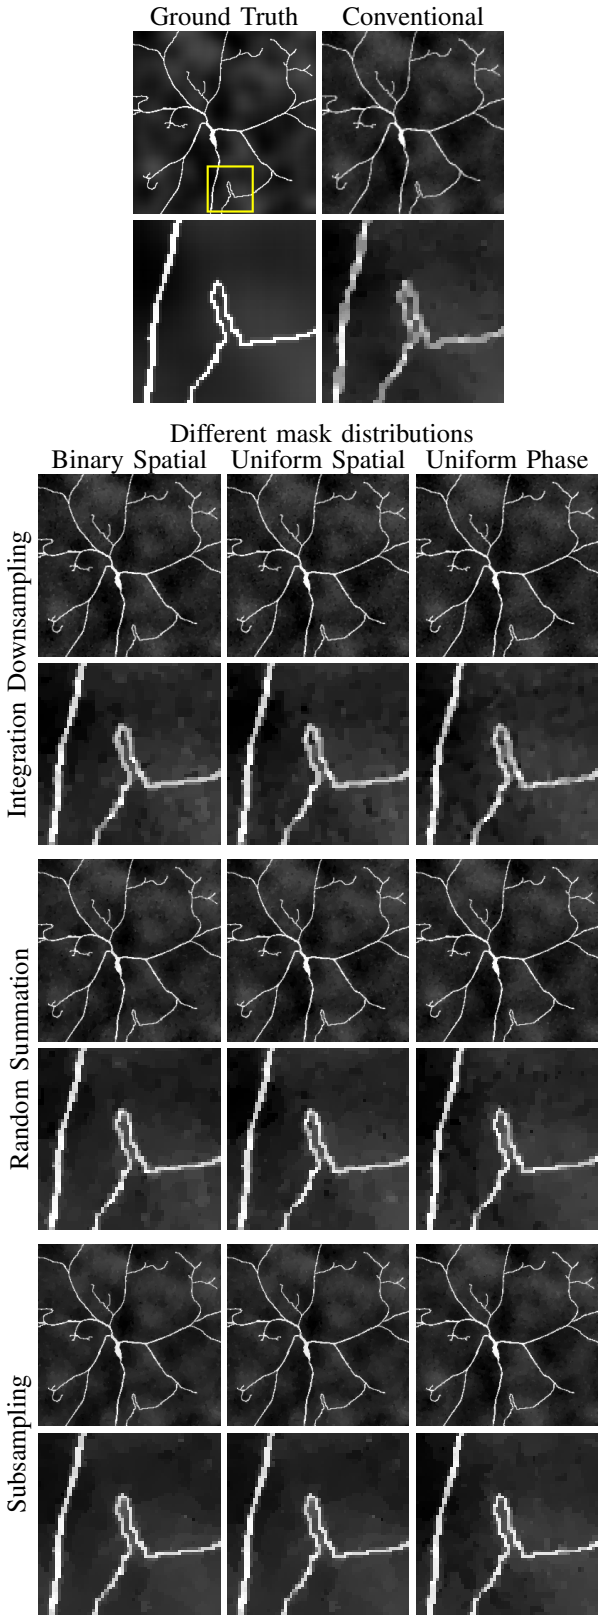

Fig. 2. Best performing reconstructions for the different downsampling strategies (integration downsampling, random summation, and subsampling) and for each downsampling strategy, the different mask distributions (binary spatial, uniform spatial, and uniform phase). For each set of images, the bottom shows the region indicated by the yellow square in the ground truth image. Clearly, significant gains in performance can be achieved by the compressive architectures over conventional imaging.

of our solution is $\mu_f = \mu_y/\mu_A n$, and we can use the zero-mean quantities in our data fidelity term

$$\phi^0(f^0) = \tfrac{1}{2}\|A^0 f^0 - y^0\|_2^2 \tag{17}$$

within the iterative algorithm, which now has the desired property that $\nabla^2\phi^0(f^0) \approx \alpha I$.

We have motivated this mean subtraction approach from an estimation setting where we first estimate the mean, then deviations about the mean. However, it can be equally motivated from a pure optimization perspective by viewing the mean subtraction step as a preconditioning strategy. The gradient-based reconstruction methods considered are known to exhibit poor convergence when the Hessian has large condition number, and the mean subtraction operation essentially improves the conditioning of the problem by eliminating the single dominant eigenvalue that occurs due to the nonzero mean.

E. Experimental Comparison on Static Images

Here we present numerical experiments supporting the proposed architectures for snapshot compressive imaging. We compare the simulated performance of the described imaging architectures in the task of high-fidelity image reconstruction. We utilize a 256×256 pixel realistic phantom image of a neuron (ground truth image in Fig. 2) which is designed to test the resolution limits of our architectures with its high-resolution features. This image is inspired by a real microscopic image found at http://www.lsi.umich.edu/facultyresearch/labs/bingye/research. We examine a traditional imaging system and a total of six convolution-based compressive imaging systems. These systems are constructed by considering each combination of the following design considerations:

- the mask generation method, including binary spatial (BS) and uniform phase (UP) models,
- switching among integration downsamping, pseudo-random summation (randomly summing $d/2$ values in each $d_1 \times d_2$ block), and subsampling.

To analyze the effect of the number of measurements we collect, we vary the scene-to-sensor downsampling ratio d to be either 4, 16, or 64 (*i.e.,* downsampling in both directions by a factor of 2, 4, or 8). While each architecture has its own unique considerations in terms of signal-to-noise efficiency (see Section III-C), we choose to normalize the experiments in such a way as to keep a constant signal-to-noise ratio at the detector. This is achieved by fixing the variance of the additive white Gaussian noise to be $\sigma^2 = \text{var}(Af)/16$ for each architecture when we choose a downsampling ratio of $d = 4$. This variance is then fixed for each architecture as we vary the downsampling ratio, the rationale being that this allows us to showcase the impact of the various subsampling operations.

We consider a comparison of many penalization schemes for image reconstruction, such as sparsity in an orthonormal wavelet (Haar) basis, isotropic and anisotropic total-variation [46], as well as an ℓ_1 sparsity penalty in the overcomplete curvelet basis [47]. Curvelets are similar to wavelets in that they capture more spatially localized information than the Fourier components, however they are designed to be more

well-adapted for capturing curvilinear singularities in images (*e.g.,* edges). To reconstruct the image, we use our own implementation of the SpaRSA algorithm [42] which utilizes the mean subtraction procedure described in Section III-D2. A coarse initialization is found by solving an unpenalized least-squares minimization via a conjugate gradient method using only the compressive data. We terminate the reconstruction when the relative change in the iterates falls below a pre-specified tolerance of 1e-3. We choose the regularization weighting to minimize the final reconstruction RMSE, measured by $\text{RMSE}(\widehat{f}) = \|\widehat{f} - f^\star\|_2/\|f^\star\|_2$. A table of the resulting RMSE values is presented in Table I.

As evidenced by both the RMSE values and the images, significant gains in performance can be achieved by the compressive architectures. The UP mask architectures slightly outperform the BS-based approaches in this simulation, but the type of subsampling has a greater effect on the quality of the reconstruction. Focusing on the $d = 4$ case, we see that subsampling performs best overall, followed by random summation, then integration downsampling. This is readily apparent from the RMSE values and the greater clarity by which we can resolve fine-scale features in the dendrites of the neuron. This is exactly predicted by the theory: larger coherence yields poorer performance. However, pushing the level of undersampling, we see that simple subsampling is not always the solution, since the performance degrades quite rapidly, as the per-measurement signal to noise ratio does not naturally increase as it does with the other subsampling architectures. An optimal architecture must strike a careful balance between low coherence and robustness to noise.

IV. Coded Aperture Keyed Exposure Sensing for Dynamic Scenes

Here we detail our compressive video acquisition method that combines coded apertures and keyed exposures to address all the competing challenges detailed in Sec. I-A. Specifically, in our CAKE imaging method, each observed frame y_k is given by an exposure of B high-rate coded observations:

$$y_k = \sum_{t=1}^{B} A_{(k-1)B+t} f^\star_{(k-1)B+t} + w_k, \tag{18}$$

where each A_t describes an AMM CCA sensing matrix. Note that since in our theory, the downsampling operator D_{sub} is a structured nonrandom operator, we can rewrite the above as

$$y_k = D_{\text{sub}} \left[\sum_{t=1}^{B} R_{(k-1)B+t} f^\star_{(k-1)B+t} \right] + w_k. \tag{19}$$

Hence all that is required to implement this sensing paradigm is modulating the coded aperture mask over the B-frame exposure time. Because of this, one can think of our system as performing a coded exposure acquisition *for each coded aperture element*.

Since our sensing strategy is independent across each low-resolution frame, and also for simplicity of presentation, we only consider the recovery of a length B block of frames from a single snapshot image in our theoretical analysis.

Theorem 4 (Coded Aperture Keyed Exposure Sensing)**.** *Let $A = [A_1 \cdots A_B]$ be an $m \times nB$ sensing matrix for the CAKE system where the coded aperture pattern for each A_t is drawn i.i.d. from a suitable probability distribution. Then there exists constants $c_1, c_2 \geq 0$ depending on δ_s such that for any*

$$m \geq c_1 s^2 \log(nB), \tag{20}$$

A is RIP(s, δ_s) with probability exceeding

$$1 - 2n^2 B^2 e^{-c_2 m/s^2}. \tag{21}$$

Note here that s denotes the sparsity of the first B frames of the video sequence, rather than simply the sparsity of an individual frame as in Thm. 2. The proof of Thm. 4 is presented in Appendix B. Similar to Thm. 2, this proof assumes that the entries of each h_t, $t = 1, \ldots, B$ are generated i.i.d. according to the scaled Rademacher distribution (9). Again, it is straightforward to extend the result to other mask generating distributions.

A. Sparse Transformation Sensing

It is most often the case that the original frames $f^\star$ are not sparse, but can be sparsely represented by a temporal transform of the frames. For simplicity, consider the first coded measurement. Within this acquisition, the coefficient sequence

$$\theta^\star_k = \sum_{t=1}^{B} W_{k,t} f^\star_t, \quad k = 1, \ldots, B,$$

may be more sparse for a well-chosen sparse temporal transform W. Notationally this is equivalent to $\theta^\star = (W \otimes I_n) f^\star$, where $\otimes$ denotes the Kronecker matrix product. A preferred sensing strategy would be to use our RIP-satisfying CAKE acquisition to sense the coefficient sequence $\theta^\star$:

$$y = \sum_{k=1}^{B} A_k \theta^\star_k + w. \tag{22}$$

Surprisingly, this can be accomplished using an identical architecture, with some slight adjustment to the coded aperture mask patterns used during the keyed acquisition.

To see this, we examine the resulting sensing matrix in terms of $f^\star$:

$$\begin{aligned} \sum_{k=1}^{B} A_k \theta^\star_k = \sum_{k=1}^{B} A_k \sum_{t=1}^{B} W_{k,t} f^\star_t &= \sum_{t=1}^{B} \left[\sum_{k=1}^{B} W_{k,t} A_k \right] f^\star_t \\ &= \sum_{t=1}^{B} A^W_t f^\star_t, \end{aligned}$$

where we define

$$A^W_t = \sum_{k=1}^{B} W_{k,t} A_k. \tag{23}$$

Therefore (22) is also a CAKE acquisition using the sensing matrices A^W_t. Because of the linear dependence on the generating mask patterns, this amounts to using *the an identical architecture with adjusted mask patterns*

$$h^W_t = \sum_{k=1}^{B} W_{k,t} h_k. \tag{24}$$

Sensing Architecture			Reconstruction RMSE (%)				
Subsampler	Mask Model	d	CG Init	ℓ_2-Haar	ℓ_2-Aniso. TV	ℓ_2-Iso. TV	ℓ_2-Curvelet
Integration downsampling	BS	4	43.552652	31.908468	**23.632185**	23.712212	34.552177
		16	52.508761	47.604093	44.602132	43.339913	43.606292
		64	60.435296	58.592028	57.537579	56.960604	56.747754
	US	4	43.220564	31.861358	**23.497272**	23.643347	34.325474
		16	52.126070	47.551547	44.400403	43.069598	43.610772
		64	59.642746	58.145647	57.118346	56.684747	56.541953
	UP	4	38.474720	30.797415	23.265628	**23.214944**	32.475710
		16	50.284641	46.653984	44.638167	43.091206	44.231886
		64	58.815932	58.168464	57.184195	56.785933	56.558017
Random Summation	BS	4	56.890247	34.101763	**18.992251**	20.535913	35.202771
		16	62.274874	51.430633	45.573628	44.644012	45.537800
		64	65.672366	60.160050	58.296549	57.789539	57.630494
	US	4	56.544494	33.659520	**17.484428**	19.021174	34.322115
		16	61.950795	50.977833	44.857052	43.641823	45.217990
		64	65.350187	59.697311	57.573941	57.159371	57.096311
	UP	4	54.473107	29.831022	**16.033758**	17.728624	33.651808
		16	60.297045	48.902002	44.939914	43.733698	45.219903
		64	64.317781	60.909768	57.360741	57.106131	56.933892
Pointwise Subsampling	BS	4	61.578677	39.742969	**17.129559**	18.611897	35.793504
		16	67.787137	65.569753	51.912236	50.842530	52.625216
		64	69.244807	69.202225	64.772033	64.693113	64.741188
	US	4	61.578677	39.780250	**17.200487**	18.703763	35.813267
		16	67.787137	65.570111	51.928327	50.883713	52.629357
		64	69.244807	69.202225	64.791854	64.687147	64.741653
	UP	4	61.005834	36.260055	**14.108310**	15.883293	34.282494
		16	67.675516	64.606384	48.280444	47.379639	48.924798
		64	69.187676	69.189077	63.591662	63.514672	63.350203
Conventional Imager		4	38.157537	36.654837	36.067380	**32.946067**	33.157436
		16	50.137698	50.117667	49.836071	47.932575	45.466410
		64	58.873484	58.873455	58.472819	57.806412	56.764936

TABLE I

SUMMARY OF THE RMSE VALUES FOR THE DIFFERENT MASK GENERATION AND DOWNSAMPLING SCHEMES, INCLUDING THE CONJUGATE GRADIENT INITIALIZATION, AND THE VARIOUS RECONSTRUCTION METHODS CONSIDERED. THE IMAGES ASSOCIATED WITH THE BEST PERFORMING METHODS INDICATED BY THE BOLD RMSE VALUES ARE SHOWN IN FIG. 2

If we denote $h^W = (h_1^W, \ldots, h_B^W)$, then this can be written more simply as $h^W = (W^T \otimes I_n)h = (W \otimes I_n)^T h$. Hence the CAKE system can adapt to any temporal sparse coding of the video sequence over the block of coded frames. In summary, to incorporate the sparse transform W applied to the frames, we simply apply the adjoint transform to the independently generated mask patterns.

A useful transformation that allows the exploitation of dependencies between frames is to assume that the *difference* between subsequent frames are sparse. In this case, we select $W = \nabla$, as defined in (12). In particular, we sense using the coded sequence of aperture patterns h^∇ where

$$h_t^\nabla = \begin{cases} h_k - h_{k+1} & k = 1, \ldots, B-1, \\ h_B & k = B, \end{cases}$$

and the generating aperture patterns h are drawn from a suitable distribution.

It is of interest to note that since the CAKE sensing matrix is a concatenation of downsampled Toeplitz matricies, this sensing strategy has clear connections to [48] where they consider concatenations of Toeplitz matricies as a sensing matrix for performing multiuser detection in wireless sensor networks. The important conceptual link is that their sensing matrix is used to determine a sparse set of simultaneously active users, where in our system, we are using it to infer a sequence of sparse frames, or a sequence of sparse difference frames.

B. Implementation and Normalization Details

Recall that these RIP-satisfying sensing matrices are generated using a zero-mean probability distribution and cannot be directly implemented in an optical system. As before, we apply an affine transformation to each of the A_k so that the resulting mask elements for each h_k are within the range $[0, 1/n]$. If we are generating a sensing matrix that satisfies RIP with respect to the difference frames, then even a binary generating distribution will cause the resulting mask patterns to have elements that are neither fully open nor fully closed. This can be accomplished using half-toning as discussed in Sec. III-C. If such an implementation is difficult, then one strategy would be to set these intermediate values to 0 or $1/n$ at random with equal probability.

C. Reconstruction for Video

Given measurements from the proposed CAKE imaging system, where we have designed the mask patterns for sparsity

Fig. 3. Example frame (frame 21) from the ground truth video sequence used in the numerical experiments.

in a temporal transform, we recover the frames by solving

$$\widehat{\theta} = \arg\min_{\theta} \tfrac{1}{2}\|A\theta - y\| + \tau\|\theta\|_1$$
$$\widehat{f} = (W^{-1} \otimes I_n)\widehat{\theta}.$$

For the case of sparse difference frames ($W = \nabla$), improvements in empirical performance are made by penalizing the total variation of the first frame, instead of simply the sparsity in that frame:

$$\widehat{\theta} = \arg\min_{\theta} \tfrac{1}{2}\|A\theta - y\| + \tau_{\text{TV}}\|\theta_1\|_{\text{TV}} + \tau_1 \sum_{k=2}^{B} \|\theta_k\|_1$$
$$\widehat{f} = (L \otimes I_n)\widehat{\theta},$$

here $L = \nabla^{-1}$ is a lower triangular matrix of all ones.

In addition, by using more than one low-rate observation per reconstruction step, we are typically able to improve performance by coding the difference frames across more than one length-B block of high-rate frames. This is the reconstruction technique we use in the experimental results section. In previous work, we noticed that when the amount of processing time allotted per frame is held constant, the accuracy generally increases with the number of frames processed simultaneously. However, one simply cannot solve for arbitrarily many frames simultaneously, as the improvement in accuracy diminishes when the size of the problem is such that only a very small number of reconstruction iterations can be run within the allotted time. We refer the reader to [26] for details.

For video sequences of longer duration, we may wish to process the video in an online fashion, solving for only a few blocks of frames simultaneously. In many applications, such as surveillance and monitoring, the video frames are strongly correlated, and the solution to a previous block of frames may be used as an initialization to the algorithm to solve the next block of frames. This estimate will generally be very accurate, and therefore, relatively few iterations are needed to obtain a solution to each optimization problem.

D. Numerical Experiments

In this section, we demonstrate the effectiveness of the proposed CAKE architecture at successfully recovering a video sequence of short-wave infrared (SWIR) data collected (courtesy of Jon Nichols at NRL) by a short-wave IR ($0.9-1.7\mu$m) camera. The camera is based on a 1024×1280 InGaAs (Indium, gallium arsenide) focal plane array with 20μm pixel pitch. Optically, the camera was built around a fixed, $f/2$ aperture and provides a 6° field of view along the diagonal with a focus range of $50\text{m} \to \infty$. Imagery were output at the standard 30Hz frame rate with a 14 bit dynamic range. An example frame is shown in Fig. 3. In this sequence, the three boats are traveling at different velocities with respect to the slowly-changing background of the waves. The size of each frame is 128×256, and we consider reconstructing 28 frames of the sequence.

We consider CAKE observations where we downsample spatially by a factor of 2 in both directions ($d_1 = d_2 = 2$), and use a coded exposure block length of $B = 4$. We compare acquiring the sequence with independent mask codes, and mask codes designed to exploit the sparse difference frames directly. We reconstruct the entire video sequence using all 7 low-resolution low-rate frames using a total variation penalty on the first frame, and ℓ_1 sparsity penalty on all subsequent difference frames. We optimize the regularization parameters to minimize the reconstruction error. For comparison, we consider traditionally captured data (*i.e.,* by simply averaging over $d_1 \times d_2 \times B$ blocks of the spatiotemporal video sequence). To interpolate this data to the original resolution of the video sequence, we consider both nearest-neighbor and spline interpolation.

We show the estimates for the different acquisition and reconstruction methods in Fig. 4. Here we focus only on a ROI of the sailboat. We see that the CAKE sensing is able to reconstruct the scene with a higher spatial and temporal resolution. This is evidenced in the residuals, which include much less critical scene structure than the conventional system. We see from the examination of the difference frame that the nearest-neighbor reconstruction from conventionally sampled data yields no motion over the two frames and suffers from poor spatial resolution. Using spline interpolation helps improve the spatial resolution, but it is still insufficient to recover the scene with high temporal resolution, as can be seen in the blur on the leading edge of the sail. Numerically we quantify the performance over the entire video sequence in terms of the RMSE (%), $100 \cdot \|\widehat{f} - f^\star\|_2 / \|f^\star\|_2$, calculated both over the entire frame, and only over the sailboat ROI. This is tabulated in Table II. In summary we see that the CAKE acquisitions are able to outperform traditionally sampled video in terms of reconstruction accuracy and reconstructing salient motion.

It should be noted that there are limitations to the CAKE system in the presence of strong motion. In this case, the sparsity level of the difference frames may drastically increase as the previous frame ceases to be a good prediction of the next frame. As such, the RIP bound in Theorem 4 states the number of measurements we require to reconstruct the scene must necessarily increase, requiring either a faster temporal resolution measurements or higher resolution FPA to achieve the same accuracy. Because of this balance, a system designer may need to make important engineering tradeoffs to implement CAKE acquisition for a particular application.

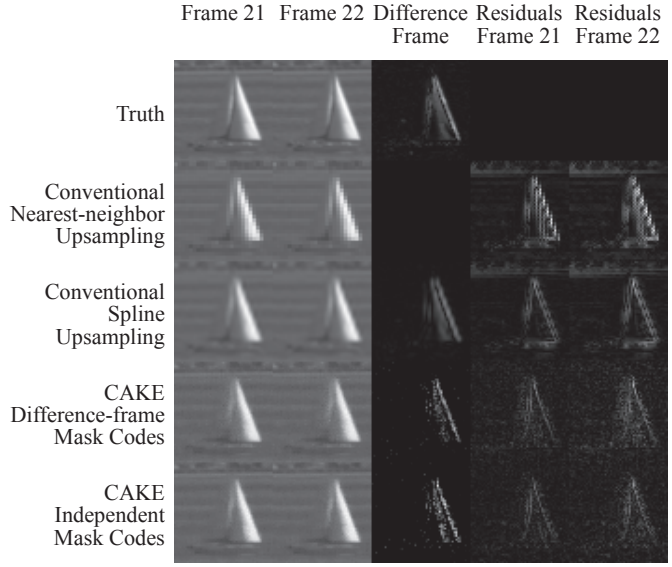

Fig. 4. Results obtained for the sailboat ROI for an example pair of frames (frames 21 and 22). Shown are the true frames, the reconstruction from traditionally sampled data, and the reconstruction from CAKE sensed data. For comparison we show the residuals as compared with the truth, as well as the difference between the frames. Note that using nearest neighbor interpolation results in no motion over the block of frame, hence a zero difference frame.

	Reconstruction RMSE (%)	
Sensing Architecture	Full Frame	Sailboat ROI
Conventional (Nearest-neighbor)	5.5163 (5.5305)	10.3525 (10.4241)
Conventional (Spline)	3.7654 (3.8047)	5.9335 (6.0377)
CAKE (Difference-frame Codes)	3.5840 (3.9183)	5.6971 (7.0421)
CAKE (Independent Codes)	**2.9079** (3.0932)	**4.3266** (4.9604)

TABLE II
RECONSTRUCTION RMSE ACHIEVED FOR THE CONVENTIONAL AND CAKE ARCHITECTURES OVER THE VIDEO SEQUENCE. RESULTS ARE REPORTED BOTH FOR THE FULL FRAME AND THE ROI OF THE SAILBOAT. DUE TO BOUNDARY ISSUES, WE REPORT THE RMSE DISCOUNTING THE FIRST AND LAST BLOCK OF B FRAMES. THE RMSE VALUES FOR THE ENTIRE SEQUENCE ARE GIVEN IN PARENTHESES.

V. CONCLUSIONS

Compressed sensing offers a strong theoretical foundation for sparse signal recovery. However, practical and implementable imaging system designs based on CS theory have lagged far behind. In this paper, we demonstrate how CS principles can be applied to physically realizable optical system designs, namely coded aperture imaging. Numerical experiments show that CS methods can help resolve high resolution features in images and videos that conventional imaging systems cannot. We have also demonstrated that our CAKE acquisition system can recover video sequences from highly under-sampled data in much broader settings than those considered in initial coded exposure studies. However, our derived theoretical limits show that there are important tradeoffs involved that depend on the spatial sparsity, and temporal similarity of the scene that ultimately govern the accuracy of our reconstructions for a specified number of compressive measurements. Finally, we note that our proposed approach performs well in high signal-to-noise ratio settings, but like all CS reconstructions, it can exhibit significant artifacts in photon-limited settings.

Two practical issues associated with coded aperture imaging in general are the blur due to misalignment of the mask and diffraction and interference effects. Noncompressive aperture codes have been developed to be robust to these effects [49]. One important avenue for future research is the development of *compressive* coded aperture makes with similar robustness properties. Noncompressive coded apertures have also been shown useful in inferring the depth of different objects in a scene; similar inference may be possible with the compressive coded apertures described in this paper.

ACKNOWLEDGEMENTS

The authors would like to thank Prof. Justin Romberg and Dr. Kerkil Choi for several valuable discussions.

APPENDIX

A. Proof of the RIP for Spatial-Domain CCA Sensing

Here we present the proof that appropriately scaled compressive coded apertures satisfy the RIP.

Proof of Theorem 2: In the interest of notational simplicity, it is easier to work on the two-dimensional images versus their one-dimensional vectorial representations. For concreteness, we assume the entries of h are generated i.i.d. according to the scaled Rademacher distribution

$$h_{k_1,k_2} = \begin{cases} \sqrt{d/n} & \text{with probability } 1/2, \\ -\sqrt{d/n} & \text{with probability } 1/2, \end{cases}$$

for $k_1 = 1, \ldots, n_1$ and $k_2 = 1, \ldots, n_2$. The proof uses the same techniques as that of Theorem 4 in [36], where the RIP is established by shifting the analysis of the submatrices of A to the entries of the Gram matrix $G = A^T A$ by invoking Geršgorin's disc theorem [50]. This theorem states that the eigenvalues of an $n \times n$ complex matrix G all lie in the union of n discs $d_j(c_j, r_j)$, $j = 1, 2, \ldots, n$, centered at $c_j = G_{j,j}$ with radius

$$r_j = \sum_{\substack{i=1 \\ i \neq j}}^{m} |G_{i,j}|.$$

In essence, we show that with high probability $G \approx I$, so that the eigenvalues of G are clustered around one with suitably high probability.

From the discussion above, we have that

$$A_{l,k} = h_{\text{mod}[(l_1-1)d_1-k_1+1,n_1]+1,\text{mod}[(l_2-1)d_2-k_2+1,n_2]+1},$$

so the entries of the resulting Gram matrix G are

$$G_{p,q} = \sum_{l_1=1}^{n_1/d_1} \sum_{l_2=1}^{n_2/d_2} h_{\text{mod}[(l_1-1)d_1-p_1+1,n_1]+1,\text{mod}[(l_2-1)d_2-p_2+1,n_2]+1} \cdot h_{\text{mod}[(l_1-1)d_1-q_1+1,n_1]+1,\text{mod}[(l_2-1)d_2-q_2+1,n_2]+1}. \quad (25)$$

From the normalization introduced for h, we can show that $\mathbb{E}[G] = I$. Now we need to bound the deviation about the mean via concentration. We consider first the diagonal entries

of G, when $p = q$. Each term is a sum of $n = n_1 n_2$ bounded i.i.d. entires, and applying Hoeffding's inequality yields

$$\mathbb{P}(|G_{q,q} - 1| \geq \delta_{\mathrm{d}}) \leq 2 \exp\left(\frac{-2n\delta_{\mathrm{d}}^2}{d}\right).$$

Next we consider the off-diagonal entries in the special case that either $\mathrm{mod}(p_1 - q_1, d_1) \neq 0$ or $\mathrm{mod}(p_2 - q_2, d_2) = 0$. In this case, each of the terms in the summand in (25) picks out a *different* set of coefficients from h, and hence there are no dependencies between different terms in the sum, hence it is also a sum of $n = n_1 n_2$ bounded i.i.d. entries, and thus from Hoeffding's inequality, we have

$$\mathbb{P}(|G_{p,q}| \geq \delta_{\mathrm{o}}/s) \leq 2 \exp\left(\frac{-n\delta_{\mathrm{o}}^2}{2ds^2}\right).$$

Lastly, we have the case that both $\mathrm{mod}(p_1 - q_1, d_1) = 0$ and $\mathrm{mod}(p_2 - q_2, d_2) = 0$ with $p \neq q$. This case deserves special care since each of the terms in the summand in (25) picks out *the same* set of coefficients from h. Therefore there necessarily exist dependencies within the terms of the sum. However, due to the special nature by which we select the coefficients, we can partition the sum $G_{p,q}$ into two sums, denoted S_1 and S_2 such that each of these is a sum of $n/2d$ independent terms. We can then apply the Hoeffding bound to each of these to yield

$$\mathbb{P}(|S_i| \geq \delta_{\mathrm{o}}/2s) \leq 2 \exp\left(\frac{-n\delta_{\mathrm{o}}^2}{4ds^2}\right), i = 1, 2.$$

Then applying the union bound gives us

$$\mathbb{P}(|G_{p,q}| \geq \delta_{\mathrm{o}}/s) \leq 4 \exp\left(\frac{-n\delta_{\mathrm{o}}^2}{4ds^2}\right).$$

Since this latter bound decays more slowly, and in the interest of simplicity, we overbound using this latter expression. What remains is to now apply the union bound over all the diagonal and off-diagonal elements. For this we need to count the number of elements in the union. For the diagonal entries, there are clearly n elements. And since the entries of G are symmetric, we only have $n(n-1)/2$ remaining terms. Hence taking $\delta_{\mathrm{d}} = \delta_{\mathrm{o}} = \delta_s/2$,

$$\begin{aligned}
&\mathbb{P}(A \text{ does not satisfy RIP}(s, \delta_s)) \\
&\quad \leq 2n \exp\left(\frac{-n\delta_s^2}{2d}\right) + 4[n(n-1)/2] \exp\left(\frac{-n\delta_s^2}{16ds^2}\right) \\
&\quad \leq [2n + 2n(n-1)] \exp\left(\frac{-n\delta_s^2}{16ds^2}\right) \\
&\quad = 2n^2 \exp\left(\frac{-n\delta_s^2}{16ds^2}\right) \leq 2n^2 \exp\left(\frac{-c_2 n}{ds^2}\right),
\end{aligned}$$

where $c_2 \leq \delta_s^2/16$. If $n \geq 2$, this probability is less than one provided $n/d \geq c_1 s^2 \log(n)$ where $c_1 \geq 3/c_2$, noting that $m = n/d$ establishes the theorem. ∎

B. Proof of RIP for CAKE

Proof of Theorem 4: This section details the proof of Theorem 4. Our strategy here is the same as that of a single frame, we bound the entries of the Gram matrix. First note that in this case the Gram matrix has a certain block structure; since $A = [A_1 \cdots A_B]$, we have

$$G = A^T A = \begin{bmatrix} A_1^T A_1 & A_1^T A_2 & \cdots & A_1^T A_B \\ A_2^T A_1 & A_2^T A_2 & \cdots & A_2^T A_B \\ \vdots & \vdots & \ddots & \vdots \\ A_B^T A_1 & A_B^T A_2 & \cdots & A_B^T A_B \end{bmatrix}.$$

This block structure allows us to utilize the results established in Appendix A. For the diagonal elements of G, we simply use the same bound for one particular block $A_k^T A_k$, since this is the situation for a sensing a single frame. Therefore

$$\mathbb{P}(|G_{q,q} - 1| \geq \delta_{\mathrm{d}}) \leq 2 \exp\left(\frac{-2n\delta_{\mathrm{d}}^2}{d}\right).$$

For the off-diagonal entries of G, we need to consider the off-diagonal entries of $A_k^T A_k$ and any entry of $A_k^T A_l$, $k \neq l$. Recall that in the Appendix A, we simply used the worst-case slower-decaying bound when the off-diagonal entry of the Gram matrix was no longer a sum of all independent terms. Here we use the same overbound and hence have

$$\mathbb{P}(|G_{p,q}| \geq \delta_{\mathrm{o}}/s) \leq 2 \exp\left(\frac{-n\delta_{\mathrm{o}}^2}{2ds^2}\right).$$

Lastly, we need to perform a union bound over all the entries of the matrix. We have nB diagonal entries, and exploiting symmetry, we only have $nB(nB-1)/2$ remaining entries. Hence taking $\delta_{\mathrm{d}} = \delta_{\mathrm{o}} = \delta_s/2$,

$$\begin{aligned}
&\mathbb{P}(A \text{ does not satisfy RIP}(s, \delta_s)) \\
&\quad \leq 2nB \exp\left(\frac{-n\delta_s^2}{2d}\right) + 4[nB(nB-1)/2] \exp\left(\frac{-n\delta_s^2}{16ds^2}\right) \\
&\quad \leq [2nB + 2nB(nB-1)] \exp\left(\frac{-n\delta_s^2}{16ds^2}\right) \\
&\quad = 2n^2B^2 \exp\left(\frac{-n\delta_s^2}{16ds^2}\right) \leq 2n^2B^2 \exp\left(\frac{-c_2 n}{ds^2}\right),
\end{aligned}$$

where $c_2 \leq \delta_s^2/16$. If $nB \geq 2$, this probability is less than one provided $n/d \geq c_1 s^2 \log(nB)$ where $c_1 \geq 3/c_2$, again noting that $m = n/d$ establishes the theorem. ∎

REFERENCES

[1] E. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inform. Theory*, vol. 52, no. 2, pp. 489 – 509, 2006.

[2] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inform. Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.

[3] M. Shankar, R. Willett, N. Pitsianis, T. Schulz, R. Gibbons, R. T. Kolste, J. Carriere, C. Chen, D. Prather, and D. Brady, "Thin infrared imaging systems through multichannel sampling," *Applied Optics*, vol. 47, no. 10, pp. B1–B10, 2008.

[4] A. F. Coskun, I. Sencan, T.-W. Su, and A. Ozcan, "Lensless wide-field flourescent imaging on a chip using compressive decoding of sparse objects," *Optics Express*, vol. 18, no. 10, pp. 10510–10523, 2010.

[5] L. C. Potter, E. Ertin, J. T. Parker, and M. Cetin, "Sparsity and compressed sensing in radar imaging," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 1006–1020, 2010.

[6] J. Gu, S. K. Nayar, E. Grinspun, P. N. Belhumeur, and R. Ramamoorthi, "Compressive structured light for recovering inhomogeneous participating media," in *Proceedings of the European Conference on Computer Vision*, 2008.

[7] M. Mohtashemi, H. Smith, D. Walburger, F. Sutton, and J. Diggans, "Sparse sensing dna microarray-based biosensor: Is it feasible?," in *2010 IEEE Sensors Applications Symposium*, Feb 2010, pp. 127–130.

[8] M.A. Sheikh, O. Milenkovic, and R.G. Baraniuk, "Designing compressive sensing dna microarrays," in *2nd IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing, 2007*, December 2007, pp. 141–144.

[9] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, December 2007.

[10] A.C. Gurbuz, J.H. McClellan, and W.R. Scott, "A compressive sensing data acquisition and imaging method for stepped frequency GPRs," *IEEE Transactions on Signal Processing*, vol. 57, no. 7, pp. 2640 – 2650, July 2009.

[11] P. Ye, J.L. Paredes, G.R. Arce, Y. Wu, C. Chen, and D.W. Prather, "Compressive confocal microscopy," in *IEEE International Conference on Acoustics, Speech and Signal Processing, 2009*, April 2009, pp. 429 –432.

[12] J. Bobin, J.-L. Starck, and R. Ottensamer, "Compressed sensing in astronomy," *Selected Topics in Signal Processing, IEEE Journal of*, vol. 2, no. 5, pp. 718 –726, October 2008.

[13] M. F. Duarte, M. A. Davenport, D. Takhar, J. N. Laska, T. Sun, K. F. Kelly, and R. G. Baraniuk, "Single pixel imaging via compressive sampling," *IEEE Sig. Proc. Mag.*, vol. 25, no. 2, pp. 83–91, 2008.

[14] J. G. Ables, "Fourier transform photography: a new method for X-ray astronomy," *Proceedings of the Astronomical Society of Australia*, vol. 1, pp. 172, 1968.

[15] R. H. Dicke, "Scatter-hole cameras for X-rays and gamma-rays," *Astrophysical Journal*, vol. 153, pp. L101–L106, 1968.

[16] E. Caroli, J. B. Stephen, G. Di Cocco, L. Natalucci, and A. Spizzichino, "Coded aperture imaging in X- and gamma-ray astronomy," *Space Science Reviews*, vol. 45, pp. 349–403, 1987.

[17] G. Skinner, "Imaging with coded-aperture masks," *Nucl. Instrum. Methods Phys. Res.*, vol. 221, pp. 33–40, Mar. 1984.

[18] R. Accorsi, F. Gasparini, and R. C. Lanza, "A coded aperture for high-resolution nuclear medicine planarimaging with a conventional anger camera: experimental results," *IEEE Trans. Nucl. Sci.*, vol. 28, pp. 2411–2417, 2001.

[19] G. R. Gindi, J. Arendt, H. H. Barrett, M. Y. Chiu, A. Ervin, C. L. Giles, M. A. Kujoory, E. L. Miller, and R. G. Simpson, "Imaging with rotating slit apertures and rotating collimators," *Medical Physics*, vol. 9, pp. 324–339, 1982.

[20] S. R. Meikle, R. R. Fulton, S. Eberl, M. Bahlbom, K.P. Wong, and M. J. Fulham, "An investigation of coded aperture imaging for small animal SPECT," *IEEE Trans. Nucl. Sci.*, vol. 28, pp. 816–821, 2001.

[21] A. Levin, R. Fergus, F. Durand, and W. T. Freeman, "Image and depth from a conventional camera with a coded aperture," in *Proc. of Intl. Conf. Comp. Graphics. and Inter. Tech.*, 2007.

[22] A. Agrawal and Yi Xu, "Coded exposure deblurring: Optimized codes for psf estimation and invertibility," in *IEEE Computer Vision and Pattern Recognition*, 2009.

[23] R. Raskar, A. Agrawal, and J. Tumblin, "Coded exposure photography: Motion deblurring using fluttered shutter," in *ACM Special Interest Group on Computer Graphics and Interactive Techniques (SIGGRAPH)*, 2006.

[24] A. Veeraraghavan, D. Reddy, and R. Raskar, "Coded strobing photography: Compressive sensing of high speed periodic videos," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, no. 4, pp. 671–686, 2011.

[25] R. F. Marcia and R. M. Willett, "Compressive coded aperture superresolution image reconstruction," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Processing (ICASSP)*, 2008.

[26] R. F. Marcia and R. M. Willett, "Compressive coded aperture video reconstruction," in *Proc. 16th Euro. Signal Process. Conf., 2008*, Lausanne, Switzerland, August 2008.

[27] R. F. Marcia, Z. T. Harmany, and R. M. Willett, "Compressive coded apertures for high-resolution imaging," in *Proc. 2010 SPIE Photonics Europe*, Brussels, Belgium, April 2010.

[28] J. Romberg, "Compressive sampling by random convolution," *SIAM Journal on Imaging Sciences*, vol. 2, no. 4, pp. 1098–1128, 2009.

[29] N. J. A. Sloane and M. Harwit, "Masks for Hadamard transform optics, and weighing designs," *Appl. Opt.*, vol. 15, no. 1, pp. 107–114, 1976.

[30] A. Ashok and M. A. Neifeld, "Pseudorandom phase masks for superresolution imaging from subpixel shifting," *Appl. Opt.*, vol. 46, no. 12, pp. 2256–2268, 2007.

[31] S. R. Gottesman and E. E. Fenimore, "New family of binary arrays for coded aperture imaging," *Appl. Opt.*, vol. 28, no. 20, pp. 4344–4352, 1989.

[32] E. J. Candès and T. Tao, "Decoding by linear programming," *IEEE Trans. Inform. Theory*, vol. 15, no. 12, pp. 4203–4215, 2005.

[33] E. Candès, "Compressive sampling," in *Int. Congress of Mathematics*, 2006, vol. 3, pp. 1433–1452.

[34] E. J. Candès, J. K. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Communications on Pure and Applied Mathematics*, vol. 59, no. 8, pp. 1207–1223, 2006.

[35] W. Bajwa, J. Haupt, G. Raz, S. Wright, and R. Nowak, "Toeplitz-structured compressed sensing matrices," in *Proc. of Stat. Sig. Proc. Workshop*, 2007.

[36] J. D. Haupt, W. U. Bajwa, G. M. Raz, and R. D. Nowak, "Toeplitz compressed sensing matrices with applications to sparse channel estimation," *IEEE Transactions on Information Theory*, vol. 56, no. 11, pp. 5862–5875, 2010.

[37] R. Robucci, J. D. Gray, L. K. Chiu, J. Romberg, and P. Hasler, "Compressive sensing on a CMOS separable-transform image sensor," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 1089–1101, 2010.

[38] V. Majidzadeh, L. Jacques, A. Schmid, P. Vandergheynst, and Y. Leblebici, "A (256*256) pixel 76.7mW CMOS imager/compressor based on real-time in-pixel compressive sensing," in *IEEE International Symposium on Circuits and Systems*, 2010, pp. 2956–2959.

[39] W. Chi and N. George, "Phase-coded aperture for optical imaging," *Optics Communications*, vol. 282, pp. 2110–2117, 2009.

[40] J. Haupt and R. Nowak, "Compressive sampling vs. conventional imaging," in *Proc. IEEE Intl. Conf. on Imaging Processing*, 2006, pp. 1269–1272.

[41] M. Raginsky, R. M. Willett, Z. T. Harmany, and R. F. Marcia, "Compressed sensing performance bounds under Poisson noise," *IEEE Trans. Signal Process.*, vol. 58, no. 8, pp. 3990–4002, 2010.

[42] S. Wright, R. Nowak, and M. Figueiredo, "Sparse reconstruction by separable approximation," *IEEE Trans. Signal Process.*, vol. 57, pp. 2479–2493, 2009.

[43] M. A. T. Figueiredo, R. D. Nowak, and S. J. Wright, "Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems," *IEEE J. Sel. Top. Sign. Proces.: Special Issue on Convex Optimization Methods for Signal Processing*, vol. 1, no. 4, pp. 586–597, 2007.

[44] J. Haupt and R. Nowak, "Signal reconstruction from noisy random projections," *IEEE Trans. Inform. Theory*, vol. 52, no. 9, pp. 4036–4048, 2006.

[45] R. G. Baraniuk, M. Davenport, R. A. DeVore, and M. B. Wakin, "A simple proof of the restricted isometry property for random matrices," *Constructive Approximation*, vol. 28, no. 3, pp. 253–263, 2008.

[46] L.I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D: Nonlinear Phenomena*, vol. 60, no. 1-4, pp. 259–268, 1992.

[47] E. J. Candès and D. L. Donoho, "New tight frames of curvelets and optimal representations of objects with piecewise C^2 singularities," *Communications on Pure and Applied Mathematics*, vol. 57, no. 2, pp. 219–266, 2004.

[48] L. Applebaum, W.U. Bajwa, M.F. Duarte, and R. Calderbank, "Asynchronous code-division random access using convex optimization," *Accepted to Elsevier Phy. Commun.*, 2011.

[49] J. W. Stayman, N. Subotic, and W. Buller, "An analysis of coded aperture acquisition and reconstruction using multi-frame code sequences for relaxed optical design constraints," in *Proceedings SPIE 7468*, 2009.

[50] R. Varga, *Geršgorin and His Circles*, Springer-Verlag, Berlin, Germany, 2004.

PLACE PHOTO HERE

Zachary T. Harmany Zachary Harmany received the B.S. (magna cum lade, with honors) in Electrical Engineering and B.S. (cum lade) in Physics from The Pennsylvania State University in 2006. Currently, he is a Ph.D. student in the department of Electrical and Computer Engineering at Duke University. In 2010 he was a visiting researcher at The University of California, Merced. His research interests include nonlinear optimization, functional neuroimaging, and signal and image processing with applications in medical imaging, astronomy, and night vision.

Roummel F. Marcia Roummel Marcia received his B.A. in Mathematics from Columbia University in 1995 and his Ph.D. in Mathematics from the University of California, San Diego in 2002. He was a Computational and Informatics in Biology and Medicine Postdoctoral Fellow in the Biochemistry Department at the University of Wisconsin-Madison and a Research Scientist in the Electrical and Computer Engineering at Duke University. He is currently an Assistant Professor of Applied Mathematics at the University of California, Merced. His research interests include nonlinear optimization, numerical linear algebra, signal and image processing, and computational biology.

Rebecca M. Willett (SM '02, M '05, SM '11) is an assistant professor in the Electrical and Computer Engineering Department at Duke University. She completed her PhD in Electrical and Computer Engineering at Rice University in 2005. Prof. Willett received the National Science Foundation CAREER Award in 2007, is a member of the DARPA Computer Science Study Group, and received an Air Force Office of Scientific Research Young Investigator Program award in 2010. Prof. Willett has also held visiting researcher positions at the Institute for Pure and Applied Mathematics at UCLA in 2004, the University of Wisconsin-Madison 2003-2005, the French National Institute for Research in Computer Science and Control (INRIA) in 2003, and the Applied Science Research and Development Laboratory at GE Healthcare in 2002. Her research interests include network and imaging science with applications in medical imaging, wireless sensor networks, astronomy, and social networks.